

ARI Research Note 90-55

DTIC FILE COPY

①

AD-A226 271

A Short Review of Human Motor Behavior: Phenomena, Theories, and Systems

Wayne Iba



University of California

for

**Contracting Officer's Representative
Judith Orasanu**

**Basic Research
Michael Kaplan, Director**

July 1990



**United States Army
Research Institute for the Behavioral and Social Sciences**

Approved for public release; distribution is unlimited.

90 09 07 012

U.S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES

**A Field Operating Agency Under the Jurisdiction
of the Deputy Chief of Staff for Personnel**

EDGAR M. JOHNSON
Technical Director

JON W. BLADES
COL, IN
Commanding

Research accomplished under contract for
the Department of the Army

University of California

Technical review by

Michael Drillings

NOTICES

DISTRIBUTION: This report has been cleared for release to the Defense Technical Information Center (DTIC) to comply with regulatory requirements. It has been given no primary distribution other than to DTIC and will be available only through DTIC or the National Technical Information Service (NTIS).

FINAL DISPOSITION: This report may be destroyed when it is no longer needed. Please do not return it to the U.S. Army Research Institute for the Behavioral and Social Sciences.

NOTE: The views, opinions, and findings in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other authorized documents.

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE


Form Approved
OMB No. 0704-0188

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS ---	
2a. SECURITY CLASSIFICATION AUTHORITY ---			3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.	
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE ---				
4. PERFORMING ORGANIZATION REPORT NUMBER(S) UCI-ICS Technical Report 89-34			5. MONITORING ORGANIZATION REPORT NUMBER(S) ARI Research Note 90-55	
6a. NAME OF PERFORMING ORGANIZATION University of California	6b. OFFICE SYMBOL (If applicable) --	7a. NAME OF MONITORING ORGANIZATION U.S. Army Research Institute for the Behavioral and Social Sciences		
6c. ADDRESS (City, State, and ZIP Code) Department of Information and Computer Science, University of California Irvine, CA 92717		7b. ADDRESS (City, State, and ZIP Code) 5001 Eisenhower Avenue Alexandria, VA 22333-5600		
8a. NAME OF FUNDING / SPONSORING ORGANIZATION U.S. Army Research Institute for the Behavioral and Social Sciences	8b. OFFICE SYMBOL (If applicable) PERI-BR	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER MDA 903-85-C-0324		
8c. ADDRESS (City, State, and ZIP Code) see 7b.		10. SOURCE OF FUNDING NUMBERS		
		PROGRAM ELEMENT NO. 61102B	PROJECT NO. 74F	TASK NO. n/a
		WORK UNIT ACCESSION NO. n/a		
11. TITLE (Include Security Classification) A Short Review of Human Motor Behavior: Phenomena, Theories, and Systems				
12. PERSONAL AUTHOR(S) Iba, Wayne				
13a. TYPE OF REPORT Interim	13b. TIME COVERED FROM 89/04 TO 89/09	14. DATE OF REPORT (Year, Month, Day) 1990, July	15. PAGE COUNT 24	
16. SUPPLEMENTARY NOTATION Judith Orasanu, contracting officer's representative				
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP		
			Machine learning, Speed-accuracy tradeoff, Motor programs	
			Motor behavior, Cognitive modeling, Robotics	
			Skill improvement, Motor schemas	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) In this paper we survey three facets of human motor behavior--phenomena, theories, and implementations. We are particularly concerned with motor behavior that exhibits improvements over time and practice; this is referred to as <u>human motor learning</u> . We begin by discussing both performance and learning phenomena that have been observed in laboratory situations. This is followed with a review of three prominent theories of human motor control from the psychological literature. The performance and learning phenomena serve as a foundation on which to compare these theories. Finally, we consider several implemented models of motor learning that have attended to constraints imposed either by the phenomena introduced earlier or by the physiological structure of the human neuromuscular system. From the material surveyed, we conclude that more research on computational models would help answer questions in human motor behavior.				
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL Judith Orasanu			22b. TELEPHONE (Include Area Code) (703) 274-8722	22c. OFFICE SYMBOL PERI-BR

ACKNOWLEDGMENT

I thank Brian Ingrassia, John Gennari, David Aha, and Pat Langley for valuable comments on an earlier draft of this paper.

WAYNE IBA



Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

A SHORT REVIEW OF HUMAN MOTOR BEHAVIOR: PHENOMENA, THEORIES, AND SYSTEMS

CONTENTS

	Page
INTRODUCTION	1
PHENOMENA OF HUMAN MOTOR CONTROL	2
PSYCHOLOGICAL THEORIES OF MOTOR CONTROL AND LEARNING	6
COMPUTATIONAL APPROACHES TO MOTOR BEHAVIOR	13
CONCLUSIONS.	16
REFERENCES	17

1. Introduction

The ability to manipulate the environment is one of the intrinsic features that demonstrates intelligence, and human intelligence is distinguished from that of other species by the sophisticated level of such manipulation. The modifications we make to our environment reflect high-level thought processes and planning; however, the basic means available for such manipulations come through the use of our arms and hands. In this respect, we are faced with the same situation as the lower animals, although human motor control takes a significant amount of time to develop relative to most animals.

Many mammals are able to walk or run within minutes of birth, whereas humans generally require about a year of development before taking their first tottering steps. Therefore, we are interested not only in how humans are able to control their limbs in interesting and skillful ways, but also in how such abilities are acquired through observation and practice. Here we see an interaction between the development of physiological and cognitive components. Physical changes are still occurring quite rapidly in the infant's brain and nervous system; such development may be related to the ability to manipulate the limbs. Likewise, an older child learning to throw bean bags at a hole clearly demonstrates cognitive processes at work, as she adjusts the speed and trajectory of her throw according to errors on previous throws. Before the acquisition of these cognitive abilities, the task was too difficult.

Researchers must address both planning and control issues in order to gain a greater understanding of how humans interact and manipulate their world and how they acquire this ability. This will involve understanding high-level thought processes and cognitive development, as well as the workings and phenomena of the muscular control system in both humans and animals. We would like to find a computational theory that cuts across both areas.

The study of limbed movement is called *kinesiology* or more simply *human motor behavior*. This field is largely a synthesis of muscular physiology and experimental psychology. Historically, the earliest notions on the subject were proposed by the fathers of modern psychology (e.g., James). When behaviorism became popular, interest in motor behavior died, as all actions were thought to be explained by stimulus-response theory. During World War II, interest in motor control was renewed in an attempt to understand the performance requirements for the types of tasks created by the military. This stage was largely influenced by cybernetics and control theory due to the feedback-driven nature of radar tracking and gunnery tasks. More recently researchers have focused on developing process-oriented theories that account for a range of phenomena pertaining to the control of limbs. Since then, more experimental work attempts to validate and falsify the predictions and explanations made by the various theories that have been proposed.

The purpose of this paper is to identify connections between theories of human motor behavior and the design and control of artificial manipulator systems. Furthermore, we want a computational model that incorporates both motor issues and cognitive issues. However before beginning on this goal, we must decide how to recognize a good theory when we have found one. We start by considering a number of the phenomena that have been identified from research in human motor control. In the first section, we describe the nature of these phenomena, the empirical evidence upon which they are based, and their respective implications for theories of human motor control. In the second section we focus on psychological theories of motor control, presenting three theories of human motor control. We rate each based upon how well they explain and account for the phenomena introduced in the first section and according to how suitable they are to a computational implementation as stated in our goals. Of course, complete coverage of the phenomena is not imperative, and we are simply looking for a semi-formal means of comparison. In the next section we consider systems for controlling artificial limbs. We consider these systems with respect to their adequacy as models of human motor learning. Finally, in our closing section, we return to our original goal - a computational theory of human motor learning dealing with complex behaviors. We conclude that the theories surveyed in this paper all fall at various distances around the mark but that none are completely satisfactory. We suggest directions that research should proceed in order to accomplish this goal.

2. Phenomena of Human Motor Control

In one sense, science is in the business of explaining and predicting phenomena. These phenomena are regularities in events that given similarly controlled situations, can be repeatedly verified by experimental techniques. For the purposes of this paper we will focus on phenomena that have already been observed and not on any predictions made by theories of motor control.

Learning always occurs in the context of a performance task, so we will also examine performance aspects of human motor control. We will consider these issues separately, first reviewing the performance phenomena and then the learning phenomena. We will concentrate on robust regularities that have been repeatedly studied and tested. In this paper we will be concerned mostly with *whether* a given theory or model can account for a particular phenomenon, and not as much with *how* such an explanation might be made. In each subsection, we will focus on describing the phenomena and the experiments associated with them, delaying discussion of explanations until the next section.

2.1. Performance Phenomena

The first two phenomena that we will consider reflect performance issues in the execution of motor skills. These are exhibited during the course of movements and do not depend upon any improvement in performance quality over time. That is, these phenomena are observable at any stage of learning to varying degrees of influence.

2.1.1. The Speed-Accuracy Tradeoff

Perhaps the most well studied and documented of all human motor behavior phenomena is the *speed-accuracy tradeoff*. This is the seemingly obvious regularity that the faster a particular skill is attempted the more difficult it is to perform the skill accurately.

Although others discussed this phenomena even earlier, Fitts (1954,1964) was possibly the first to rigorously examine, study, test, and report the phenomena. Fitts' careful studies led to the formulation of a relation, known as *Fitts' Law*, that captures the maxim "haste makes waste" with quantitative values. This law relates the movement time (*MT*) to the index of difficulty (*ID*),

$$MT = a + bID \quad (1)$$

That is, if the constants *a* and *b* are known (for a particular arm) then the *MT* of the arm for a task with a particular *ID* can be predicted.

Fitts (1964) motivated the index of difficulty using information theory, defining it with the equation

$$ID = \log_2 \frac{2A}{W} \quad (2)$$

This amounts to the ratio of the movement amplitude (*A*) to the target width (*W*). Now let us examine how this is demonstrated and observed in movements in the laboratory.

Fitts and Peterson (1964) manipulated two independent variables in a discrete motor task: the distance or amplitude to be moved and the width of the target to be touched. The distance of the movements were either 3, 6, or 12 inches while the widths of the targets were 1/8, 1/4, 1/2 and 1 inch. This led to 12 experimental conditions combining movement amplitude and target width.

Subjects were required to make rapid, ballistic movements to one of a pair of targets; the appropriate target was indicated by a stimulus light. The targets were replaceable with variable widths and at different distances from the starting button. The subjects would hold a stylus on the starting button and move the stylus to the appropriate target as rapidly as possible. Fitts and Peterson reported several slight variations on this procedure, but the results were essentially identical and the results conformed to the predictions made by Fitts' Law.

This law captures the complementary nature of distance and precision. It explains why writing one's name on a paper and on a black board requires comparable amounts of time; while the distance traveled on the black board has increased, the local accuracy has decreased. Therefore, the ratio of distance to accuracy remains constant, as does the movement time.

However, speed reflects movement time, and above we claimed that speed is traded off with accuracy. If so, we would not expect to deal with constant movement times, as these are all part of the same equation. If we hold distance constant and combine equations 1 and 2, we get a relation more like the speed accuracy tradeoff:

$$W = 2 \frac{A}{2^{MT-2}} \quad (3)$$

In this case, W corresponds to the expected error, and we can use this equation to predict the pattern of errors for a given length movement over variable performance speeds. By looking at this equation, one can see that it predicts that the error will increase exponentially as the speed is increased (MT is decreased) linearly.

We have claimed that the phenomena discussed in this paper are robust and well documented. This is especially the case with the speed-accuracy tradeoff. Many other studies have shown that Fitts' Law generalizes to other types of movements and also to movements using joints other than the shoulder/elbow. Langolf, Chaffin, and Foulke (1976) have demonstrated that movements of the finger, wrist, and arm all conform to Fitts' Law, but that the constants differ from one set of joints to another. That is, the wrist is more accurate than the arm and the fingers are more accurate than the wrist. These results are for finger movements of around 1/10 inch in length and wrist movements of 1/2 inch in length performed under the magnification of a microscope. Fitts' law has received considerable support and practically no evidence indicates that it may not hold.

2.1.2. Inter-limb Similarities for Skills

The other performance phenomenon that we will consider involves the similarities observed when a skill is performed on different limbs. This can be thought of as transfer of skill between limbs. More specifically, characteristics of skills learned with one limb are evident when the same skill is performed by another limb. We must be careful not to confuse this phenomena with the more widely studied issue of transfer of *learning* between *tasks* (see Schmidt, 1975a).

For example, consider a comparison of samples from someone's handwriting or signature with various limbs: preferred hand, opposite hand, foot, mouth, etc. This is a well-known demonstration, and the comparison is usually done qualitatively by simply looking at the handwriting samples and noting common characteristics (Raibert, 1976). Figure 1 shows several samples of handwriting generated by a single subject using different modalities. Ideally, however, it would be preferable to quantitatively compare samples by recording velocities and accelerations over time and comparing the oscillation patterns.

There is additional evidence for this phenomenon in Rosenbaum's (1977) study of fatigue in the *rotor task*. His study examined two basic conditions. Rosenbaum had subjects either crank a handle in a circular motion as rapidly as possible for 30 seconds, or twisted a handle back and forth for 30 seconds. With minimal interruption, the subjects were then required to crank or twist (a 2×2 factorial design) with the *other* hand as rapidly as possible. The dependent measure of interest was the speed of cranking or twisting with the second hand. The results indicated that fatigue from one task transferred to the same task but not to the other task. Although this does not exactly represent the transfer of skill between limbs, it does lend evidence that something at a higher level than the arm muscles and nerves is common among limbs.

The transfer of skills between limbs is not as well documented as the speed-accuracy tradeoff. However, together they provide a starting place from which to compare motor control models along performance dimensions. Next we consider several learning phenomena in turn.

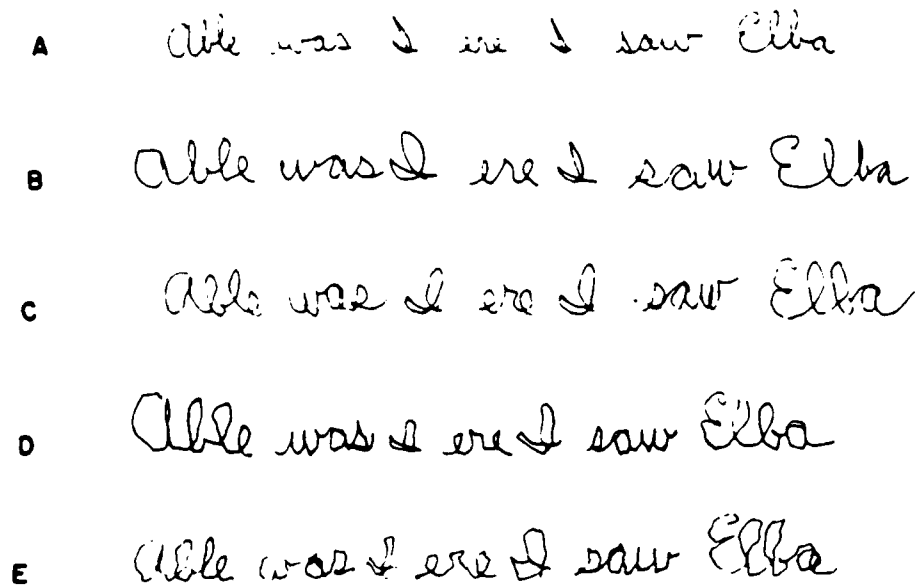


Figure 1. Five samples of handwriting from the same person using the right hand (A), right arm (B), left hand (C), mouth (D), and right foot (E) (from Raibert, 1976).

2.2. Learning Phenomena

Learning is reflected in the demonstrated improvement in performance for a particular task as a result of experience or practice. The phenomena we will consider here relate to various factors that influence the rate at which such gains in performance are acquired or the conditions under which such improvement is facilitated. Also, we will consider changes in the nature of performance as a result of learning.

2.2.1. Power Law of Practice

In general, performance appears to improve with practice, but this is not the full story. The type, quality, quantity, and scheduling of practice are all significant factors that influence to what degree (if any) improvements are gained. In this section we consider a quantitative result that relates the improvement in performance speed to the amount of practice.

This relationship has been known as the *log-log linear learning law* (Snoddy, 1926), as *DeJong's Law* (Crossman, 1959), and simply as the *power law of practice* (Newell & Rosenbloom, 1981). All versions of this law make the same claim - that a logarithmic improvement in performance speed requires a logarithmic amount of practice. The phenomenon has yet again been referred to as the *law of diminishing returns*, referring to the fact that the amount of practice necessary to improve performance increases over time.

This regularity was well documented in a study by Crossman (1959), who studied a number of workers making cigars. The cigars were made on a machine that was operated by the workers in the study. Over a period of seven years, data was collected for the same workers on how fast they were able to make a cigar.

Figure 2 shows a graph of the time to make a single cigar as a function of the number of cigars previously made. The results indicated that decreases in the time to make a cigar were achieved only after increasingly greater amounts of practice. That is, the *rate* of improvement declines with increasing practice. When plotted using log scales for the horizontal and vertical axis, the data points describe a straight line up to two years. At two years the operators appear to have stopped improving. This is attributed to the minimum




Figure 2. Cigar manufacture time as a function of the number of previous cigars manufactured on logarithmic scales (from Crossman, 1959).

cycle time of the cigar making machines; that is, after two years the operators were producing cigars in the minimum time allowed by the machinery.

Newell and Rosenbloom (1981) present a comprehensive discussion of power laws and how the experimental data fit these theoretical curves. As they point out, it is not exactly clear if the data fit closer to a power-law or to an exponential curve. They suggest that there may be other learning processes involved which mask the power-law curves. Whether it is a power law or exponential, this quantitative relation has only been demonstrated to hold for speed of performance. We might also expect it to apply to other aspects of performance, such as the amount of error and the need for attention. Although the speed and amount of error is related by the speed-accuracy tradeoff discussed above, in these types of learning studies, error is kept constant at a minimum level. Whether this relation also holds for skills such as free-throw accuracy, remains to be demonstrated. Next we turn to the need for attention during the performance of a task and how that need changes as a result of practice.

2.2.2. Transfer from Closed-loop to Open-loop Behavior

Considerable attention has been paid to the automation of skills. However, much of the discussion generated around this issue has focused on defining and identifying automation. That is, what does it mean for a skill to become "automatic" and when does such a transition occur? We will consider a trend toward automation to be a reduction in the attentional resources necessary to perform a particular task. Unfortunately, this only pushes the problem back one level. What do we mean by attention and how do we measure it? Again, many studies have been devoted to this question, but we will simply describe attention here as an emergent phenomenon. For our purposes, the amount of attention necessary for a given task is directly related to the amount of interference (in performance) caused by a coincident distraction task.

A common method of exploring this interference has been the use of a secondary reaction time task. That is, during the performance of a main motor task, the subject is required to respond to a probe as quickly as possible. The degree to which the tasks interfere should be reflected in an increased reaction time to the probe. Ellis (1969) used just such a design with a main task of moving a pointer to a target as quickly as possible and varying the temporal presentation of the probe. The results indicated that, with practice, subjects reduced their reaction times on the secondary probe task. Similar results have been found by Salmoni (1973).

Unfortunately, the results from these experiments (and many others like them) do not tell us clearly what is actually happening with respect to automation and attention. Currently there is considerable debate about the nature of attention and about skills that are said to be "automatic". Other studies have shown that combining two tasks or skills can result in interference, whereas one of the two paired with yet another task will yield no interference. For now, however, our main concern is satisfied by these results. They indicate that when two tasks *do* interfere, practice tends to reduce such interference.

This aspect of the phenomena is also closely associated with what can be called the shift from closed-loop to open-loop control (Pew, 1966). Closed-loop control implies feedback, error detection, and error correction; a movement performed in open-loop control receives no feedback and is run to completion without opportunity for adjustments. Here, the issue is the presence and use of feedback instead of the availability of attentional resources. But clearly these are closely related in so far as it requires attention to evaluate feedback information and determine what to do to improve the movement. A restatement of our phenomenon then would be that through learning a subject is able to shift motor control from a jerky, feedback-centered performance to a smooth execution of feedback-free movement.

2.2.3. Practice Variability Effects

Most of the phenomena in our list have historically been explored in their own right and then later included and explained in a particular theory of motor learning or control. The practice variability effect is unusual in this respect in that this phenomenon was predicted by Schmidt's schema theory (1975b).

The prediction made can be stated as follows: the more varied the practice, the more accurately a novel but related task will be performed. McCracken and Stelmach (1977) tested this prediction in an experiment requiring subjects to make timed movements of 200 msec. The goal was to reach a barrier marking the end of the movement distance as close to 200 msec. as possible. The length of the movement was manipulated according to the experimental conditions. There were two training conditions - high and low variability. In the high-variable condition, subjects were trained on four different length movements. In the low-variable condition, subjects were trained only on a single length movement. After training, both groups were required to perform a novel movement, where the length had not been previously performed, again in a 200 msec. time period.

The results demonstrated a weak support for the initial prediction - that the high-variable practice group would perform better on the transfer task. Although the low-variable group appear to have lower errors than the high-variable group on the initial task, the high-variable group had significantly lower errors on the transfer task. Other researchers have demonstrated similar results, and Frohlich and Elliott (1984) have extended these results beyond motor control. They have obtained variable practice effects in operating dynamic systems that are external to the human motor system. Unfortunately, there are also studies that fail to support this phenomenon (Melville, 1976) or that even present contradictory evidence (Zelaznick, 1977). Although some controversy exists around this phenomenon, it is clearly in operation in some circumstances and the question becomes one of qualifying those contexts. Therefore, a good model of human motor control should be able explain the phenomenon in some situations. Now let us turn to some of the psychological motor theories that have been proposed and see whether they can account for this phenomenon and those discussed above.

3. Psychological Theories of Motor Control and Learning

As we stated above, early motor behavior research was characterized by the identification of phenomena. Of course, this is an important stage of any developing discipline. Ultimately, however, such phenomena must be collected into a coherent story, or theory, that explains as many of the known phenomena as possible and makes predictions about new phenomena that can be verified. As predictions made by one theory are falsified, new theories arise that make the "correct" prediction and additionally make new predictions needing verification. Such is the progression of science.

This is precisely what has happened in the field of human motor behavior. Adams (1971) proposed perhaps the first comprehensive theory of human motor behavior. Almost coincidentally, Pew (1974, 1970) suggested an alternative theory that emphasized different aspects of the complete story. In response to these (and other accounts), Schmidt (1975b) proposed his own theory, which has gained acceptance and has stood the test of time quite well up to the present.

Certainly there were other theoretical results before, during, and after this period, and we are not intending to exclude this work. However, we are considering a theory to be *comprehensive* if it includes at least the following: a reasonably detailed description of the memory structures required, a detailed outline of the modules responsible for the production of motor behavior, and a careful description of the processes involved in acquiring the representations in memory used to generate movement. As an example, in this light Saltzman (1979) would not be considered as comprehensive as those mentioned above. Although he provides an extremely detailed analysis of representation structures, he only alludes to the production and acquisition components. Thus, we will consider only the theories we have mentioned above and focus on their memory structures, performance mechanisms, and learning processes.

3.1. Adams' Closed-loop Theory of Motor Learning

The scope of Adams' (1971) theory is intended to include "the instrumental learning of simple, self-paced, graded movements, like drawing a line, even though the implications extend further. And the bounds include only learning by humans old enough to have a verbal capability" (p. 122). As the title of the theory implies, it is a closed-loop, feedback-centered approach. Drawing upon early servo-mechanism ideas, Adams' model resembles the classic closed-loop control mechanism found in control theory.

3.1.1. Memory Structures

There are two basic memory structures in Adams' theory – the *perceptual trace* and the *memory trace*. The perceptual trace is memory of previous experience in movements, and the memory trace is the pattern used for generating movements.

The perceptual trace is based upon multiple sources of sensory feedback. Proprioception is a predominant source, but visual and tactual information are also very important. Even auditory feedback can be useful in many situations. For example, the sound of the ball on a bat resulting from a "good" hit is distinctive and will provide cues for predicting the result. Although the perceptual trace is thought of as a single memory structure, Adams (1971, p. 125) states that "in actuality it is a complex distribution of traces." The movement on any given trial creates a trace which contributes to the total distribution of traces. Each individual trace will tend to fade and ultimately be forgotten, but the distribution somehow manages to get stronger, although this process is not explained. The strength of the perceptual trace, thought of as a unit, is an increasing function of the number of trials on which feedback was given. As similar traces are repeated over and over, the mode of the distribution becomes strong and allows a distinctive trace to arise as the means of comparison. The perceptual trace comes to correspond to the sensations associated with the correct end point of a particular movement.

In the context of simple, self-paced movements and feedback control, the extent of a movement is the predominant controlling property. In such movements, feedback plays an integral role, but the feedback must be compared to some standard of reference to determine the correct extent of the movement. The perceptual trace performs this role in Adams' theory.

It might seem that the perceptual trace alone is sufficient for the generation and control of movement; however, there are several problems associated with this position. First, every movement will appear to be correct if it is initiated by the same structure as is used for the reference in a typical closed-loop system. Also, using only the perceptual trace as the reference of correctness, requires feedback, which is not available until approximately 200 msec. into the movement. Finally, results from verbal behavior indicate that recall and recognition, or the production and recognition of responses respectively, are based on two different memory

states (Adams & Bray, 1970; Kintsch, 1970). To account for these, Adams includes in his theory another structure called the *memory trace*.

The memory trace is introduced to "select and initiate the response, preceding the use of the perceptual trace" (p. 125). This structure is responsible for controlling a movement once initiated, until sensory feedback can be compared with the perceptual trace. The remainder of the movement is governed by feedback and the perceptual trace. Adams admits that he is uncomfortable with this form of two-state memory, but sees it as the most reasonable choice given the closed-loop assumptions and the nature of the proposed perceptual trace. He contrasts the perceptual trace, which controls the extent of a movement, with the memory trace, which controls the selection of a movement. Here the limiting context of self-paced straight line movements mentioned above is particularly evident, as more complex movements cannot be described by duration or length.

3.1.2. Producing and Improving Movements

In Adams' theory, the performance component is quite simplistic, so we will consider both performance and learning issues together. Consider how the memory structures described above are utilized to produce voluntary movements. The production of movements in Adams' theory involves using the perceptual and memory traces in a typical closed-loop feedback control system. The memory trace is the (initial) generator and selects the path to be followed. After the initial delay, feedback becomes available and the perceptual trace comes into action, controlling the remainder of the movement. The perceptual trace is compared with the sensory feedback, and adjustments are made in an effort to reach a zero error end state.

In order to improve performance, one or both of the memory structures, used to control movement, must somehow be modified. The memory trace is strengthened as a function of knowledge of results and practice. However, Adams claims that this is not the source of significant improvement. Instead, the building and strengthening of the perceptual trace is credited with improvements.

As stated above, the strength of the perceptual trace is a function of the sensory feedback experienced on each trial. Improvements could be gained simply from the drift in the mode of the distribution of sensory traces as a result of more correct sensory experience, but this implies a conscious change in the tendency of the movements. Learning actually occurs when the subject uses the knowledge of results to make the next response be different than the previous one. That is, the perceptual trace is modified and applied with respect to the previous knowledge of results.

Since movement in Adams' theory is explicitly controlled by the perceptual trace, an "average" over many similar experiences, Adams cannot explain the generation of movements that are similar, except with individual traces. This requires a separate trace for every movement ever produced, introducing a massive memory load. Below, we see that Pew (1974) presents a theory that addresses this issue, by including a more general memory structure.

3.2. Pew's Closed-loop Theory

Pew (1974) presents a closed-loop theory of human motor performance that is very similar to Adams' but with a somewhat different flavor. Although the theory is oriented towards performance issues, Pew does outline what would be involved in the acquisition of motor skills within his framework. Most of the attention is focused on performance, leaving representational issues more sketchy than in Adams' theory.

3.2.1. Memory Structures

The basic motor memory structure in Pew's theory is the *movement pattern*. This is similar to the concept of a *motor program*, in so far as it is a string of motor commands that can accept parameters to slightly alter the resulting movement along certain dimensions. The movement pattern "may be thought of as a stored representation of a path in space through which the members of the body will move" (Pew, 1974, p. 31). These patterns are stored or collected under the second memory structure - the *schema*. The idea for

schema learning is credited to Bartlett (1958) and Posner and Keele (1968), but probably goes much further back than that. However, in Pew's theory, the exact nature of the schema is even more unclear than for the movement patterns. "What properties of a movement pattern are encoded? What properties are intrinsic to a particular schema and what properties are only dimensional parameters that are free to vary from one execution to another?" (p. 28) are all questions that Pew asks but leaves unanswered.

The schema and the *schema instance* (which is nothing more than the movement pattern generated or selected from a given schema) are the necessary memory structures for the generation of movements. But as we saw in Adams' theory, this is not sufficient for the closed-loop control of voluntary movements. Pew posits that the result of selecting a particular movement pattern, the schema instance, is the generation of an image of the sensory consequences experienced when actually executing the movement pattern. The sensory consequences are analogous and perform the same role as the perceptual trace in Adams' theory. It is the image of the sensory consequences that allows the detection and correction of errors in movements while they are in progress.

3.2.2. Producing Movements

Since both Pew and Adams' present closed-loop theories, the means of movement generation will be very similar, though the memory structures used are different. In Pew's theory, a particular movement pattern is selected from the schema (the generalized source of movement information) according to the stimulating conditions existing in the environment. Of course, the selection process depends upon both the dynamic state of the subject and the environment at the current time. Once the schema instance has been selected, it must be translated into a temporal string of motor commands recognizable by the limb effectors. Pew suggests that at this stage the timing (or speed) information is added to the string of muscle commands. This allows the movement to be speeded up or slowed down as a whole. Schmidt et al. (1985), Schmidt (1982), and Armstrong (1970) present evidence that practiced movements maintain their temporal relationships independent of performance speed. This suggests a speed parameter applied to a string of motor commands that stretches and shrinks the entire movement uniformly.

Once the temporal sequence of muscle commands is formulated, all that remains is to execute this program. The muscles are then activated according to this sequence, producing a movement in space and time. However, for various reasons movements do not always proceed exactly as intended. In these cases, one needs some correction mechanism.

One interesting point about Pew's theory is that he stresses multiple levels of feedback and expected consequences. For example, he describes knowledge of results as a high-level feedback and details about the goal to be achieved as high-level expected consequences. At a lower level, the actual sensory consequences received from executing the movement pattern can be compared with the perceptual trace of expected sensory consequences. He lists these two levels as examples of a possible larger set of levels that interact during the performance of movements. Therefore, it is difficult for Pew to explicate the comparison process that results in alterations to the ongoing movement.

However, a unique point in this matter is that, in Pew's opinion, "corrections are executed ... not on the basis of deviations from a predetermined path but rather on the basis of revised estimates of where the target is with respect to where the subject's hand now is" (p. 25). This implies not only a significantly different comparison and correction mechanism from Adams', but also a much more complex one. Information from the high-level goals, the sensory consequences, and the limbs must all be integrated to allow modifications to either the schema instance selector or the actual generalized schema. Given sufficient execution time, Pew allows modifications to ongoing movements either by low-level corrective mechanisms to the movement pattern, or the initiation of a modified schema instance. But with respect to acquiring the schema structure, we are concerned with modifications that arise from previous results and how such modifications relate to changes in the same movement sometime in the future.

Pew hedges at this point and claims that, at the time of his theory, it was too early to determine the nature of the changes resulting from experience. He hazards the guess that learning involves modifications to the generalized schema structure, to the process of choosing a schema instance based upon environmental conditions, and to the nature of the implementation of the motor command sequence as generated by the movement pattern. These latter two imply that learning involves changes in the processes that control the generation of movement. In general, this is an undesirable position unless satisfactory constraints are imposed on the allowable changes. However, remember that Pew was mainly focusing on performance. Pew does make an important point about learning, once again relating to the multiple levels of feedback. He claims that the knowledge of results for a given movement is not sufficient to allow the subject to improve performance. According to Pew's model, "information about the expected sensory consequences, and about the actual sensory consequences together with the success or failure of the movement pattern, all converge in the Comparator Mechanism to produce the basis for modifications to the generalized schema, the instance selection rules, and the temporal implementation of the command sequence" (p. 32).

This broader view of feedback and comparisons, which incorporates multiple levels of information, gives Pew's theory more explanatory power than Adams' account. But before comparing these two theories, we turn to the third theory we will consider, Schmidt's schema theory, which synthesizes those of Adams and Pew.

3.3. Schmidt's Schema Theory

Adams' and Pew's theories, proposed in 1971 and 1974, spurred a flurry of experimental studies testing the predictions and claims contained therein. Schmidt proposed his schema theory (1975b) largely in response to explanatory weaknesses that were revealed as a result of these studies. However, Schmidt credits both Adams and Pew for his conceptual foundations, and the similarities to both are striking.

3.3.1. Memory Structures

Schmidt takes the ideas of the motor program (movement pattern) and the schema from Pew and develops them more fully. Pew (1974) avoided the term motor program although he did think of his schema instance as "a computer program waiting to be read" (p. 31). The motor program here is analogous to Pew's schema instance, but perhaps a bit more generalized. It is presented as requiring multiple parameters for full instantiation. Parameters include speed, as with Pew's schema instance, but also force, distance, and the possibility of others that are unmentioned. The motor program is intended to provide the means of producing a whole class of similar movements from a single memory structure. This occurs in the same way that a program designed to calculate the average of a set of numbers is usually not limited to the calculation of a single average for a fixed set of numbers. Instead, it can calculate virtually any average given the input data. In this way, Schmidt's motor program is actually a means of producing a sequence of muscle commands based upon parameters and is not the actual sequence of commands itself. The motor programs are stored collectively under, or at least are indexed through, the motor schemas.

As mentioned above, the idea of the motor schema is not new. In Schmidt's theory, it is thought of as a general rule that can be used for generating, or selecting, a motor program. In this respect it is like Pew's schema, which bundled the movement patterns. However, Schmidt proposes two different types of motor schemas – the *recall schema* and the *recognition schema* – and goes into greater detail of description than Pew. Like the work on verbal behavior and memory, the recall schema is responsible for producing movements, whereas the recognition schema is responsible for recognizing particular movements.

The recall schema is an abstraction of previous attempts at a particular class of movements. Specifically, the abstracted information includes the initial conditions at the beginning of the movement, the response specifications, and the response outcome from each movement. The initial conditions are simply a representation of the beginning state of the subject and the environment. The response specifications correspond to the parameter values used in the motor program that generated a particular movement instance. Finally, the

response outcome is a qualitative assessment of whether or not the original higher level goal was satisfied. This is commonly referred to as *knowledge of results* since there is an implied ability to make a judgement about the success of the movement. These three pieces of information are collected and stored, as in a vector, and it is the relationship among all of them that is captured as a rule or recall schema.

The recognition schema is similar to the recall schema, but instead of storing the response specifications, it stores the actual sensory consequences. As before, the sensory consequences are the trace of feedback (not limited to proprioceptive) resulting from a particular movement. Thus, the initial conditions and the response outcome are again stored, along with the sensory consequences, and the relationship among these three is abstracted to form a schema, or rule.

Finally, the error labeling schema takes the raw sensory signals coming from the limbs and the environment, and converts this input into a qualitative evaluation of the completed or ongoing movement. This labeled error signal is known as *subjective reinforcement* and can be substituted for true knowledge of results, although it will be less accurate. The error schema stores the past sensory signals along with the actual knowledge of results and builds up a rule that relates knowledge of results to the sensory signals received. Once this rule is well developed from previous experience it can be used to predict the movement outcome just from the sensory consequences.

In summary, Schmidt proposes three types of schemas - the recall, recognition, and error labeling schemas - in addition to the motor program. Next we look at how these structures are used together to produce skilled controlled movements.

3.3.2. Producing Movements

The performance component of Schmidt's theory can be split into two parts or phases - the movement preparation stage and the actual movement generation. These happen in sequence, but they can loop as well. The present theory assumes that a motor response schema (combined recall and recognition schemas) already exists.

The movement preparation stage involves taking the specified desired outcome and determining the initial conditions. Based upon the relationship developed over previous movement experience between these two variables and response specifications, the motor program is supplied with a new set of response specifications (hopefully appropriate to the situation and desired outcome). The initial conditions and desired outcome may never have been encountered before, and the resulting response specifications will be determined by "interpolating among past specifications" (p. 236). This may result in novel behaviors that have never been performed before. Simultaneously, the response schema selects the expected proprioceptive and exteroceptive feedback based upon the relationship between previous outcomes, initial conditions, and sensory consequences. Once the motor program and expected sensory consequences have been prepared, the actual movement can be initiated by running the motor program on the limb effectors.

As the muscles are activated by the motor program, the movement proceeds uninterrupted for the first 200 msec. That is, the motor program completely specifies the movement for at least this initial period. When sensory feedback becomes available, it is compared against the expected sensory consequences as given in the recognition schema. Note that the actual sensory information is coming both from the limbs and the environment, and that the expected sensory consequences likewise include multiple modalities. This comparison leads to a raw error signal which is fed back to the schemas so that adjustments may be made if necessary. The error signal is also input to the error-labeling schema for a qualitative evaluation that results in subjective reinforcement.

Once the raw error signals and subjective reinforcement are available, the entire process begins again. The desired outcome will be the same, but there will be new initial conditions and a potentially different motor response schema based upon the immediately prior movement. Each segment is performed in open-loop mode. This cycle repeats, effectively yielding closed-loop control, until the resulting error signals indicate

no further movement is necessary, or until the subjective reinforcement predicts the accomplishment of the desired outcome.

3.3.3. Modifying the Response Schemas

Schmidt proposes that the schema structures are modified by the trace from each movement. A trace starts with the initial conditions and response specifications, with the sensory consequences being added when they become available. Finally, at the end of the movement, the outcome of the movement is added to the trace, either in the form of knowledge of results or as subjective reinforcement. These four items are used to revise the means of predicting sensory consequences and response specifications on future trials. A trace is hypothesized to be rather short-lived in duration. Although this trace is unstable as a memory structure, it persists long enough to modify the recall and recognition schemas in memory.

The schemas are much more permanent memory structures that are generally resistant to forgetting. The strength of the schema increases in proportion to the number of trials of a particular class that are "sufficiently similar" to be grouped together. Also, the reliability of the relationship given in the schema increases with better quality feedback from the response outcomes.

However, the nature of the modification to the schemas is difficult to assess. Schmidt uses the term "abstraction" to describe the process of bundling up the four pieces of information described above. He states that "it is the relationship among the arrays of information that is abstracted rather than the commonalities among the elements of a single array" (p. 235). By this he seems to mean that the multi-way relationships between the four items is more important than the relationship between any particular set of initial and final conditions, response specifications, and sensory consequences. This is important because the methods for choosing the response specifications (and sensory consequences) rely on interpolating between previous experiences or using a function that is based on an interpolation of previous experiences. Recall and recognition schemas are both treated similarly with respect to learning.

The formation and modification of the error-labeling schema is even less well formulated than with the recall and recognition schemas. The strength of this schema again depends on the amount and the quality of prior experience. Previous raw error signals (the discrepancies between the expected and actual sensory states) have been stored in association with the resulting qualitative feedback (knowledge of results). Of course, the schema as a whole would have to be associated with the recall and recognition schemas to allow retrieval, since the initial and final conditions are *not* part of this memory structure. Again, as in Adams' and Pew, we see that Schmidt's theory leaves much of the learning processes to the readers' imagination. However, we can still compare these theories' learning components, their explanatory powers, and their complexities.

3.4. Analysis of the Three Theories

In this section we compare the three theories we have introduced above, with respect to their representations, performance processes, and learning methods. Although there are many similarities between these theories, they each have strengths in different aspects. All three of these psychological theories contain feedback components but only the first two, Adams and Pew's, should be considered as closed-loop theories of motor control. In these models, once the movement is going, the control is based on feedback compared with the standard of correct movement. Schmidt's theory, on the other hand, uses feedback to revise the selection of open-loop movements in the course of trying to satisfy the desired behavior designated to the motor system. In Schmidt's theory, each individual segment is considered to be under open-loop control. This actually creates a blur in the distinction between closed-loop and open-loop control.

Furthermore, Adams and Pew's theories are very much alike in form and process (with the exception of the learning processes lacking in Pew) but mainly different in representation. Adams recognizes the need for two memory structures, whereas Pew blurs this point by generating a second structure, the expected sensory consequences, from the movement pattern used to generate the movement. On the other hand, Pew's more

general memory structures allow greater flexibility in movement generation. Schmidt's overall framework bears many similarities to Pew's in representational structure, but borrows from Adams' in processes for learning and the second memory structure. From a purely theoretical and structural view, Schmidt borrows heavily from previous work but his synthesis stands as a significant improvement.

As we stated at the beginning of the paper, the purpose of considering the human phenomena was to evaluate and constrain theories of human motor learning. All of these theories can account for the speed-accuracy tradeoff by the greater number of chances to correct errors during slower movements. In addition, Schmidt (1985) presents the impulse variability theory as an alternative explanation for this phenomenon that is independent of his schema theory. This explanation is also independent of Adams' and Pew's theories and therefore, could apply in conjunction with either of these as well. However, whether the quantitative results from these theories would correspond to those predicted by Fitts' law is an open question. Such verification would require instantiating these theories as computational models - which has not yet been done. Likewise, the transfer of skills between limbs could probably be handled by appropriately transforming the memory representation for a given skill to be executed on another limb.

Since Pew's theory does not explicitly address learning issues, we cannot say much about his theory with respect to the learning phenomena. Certainly, all three theories predict improvement based upon experience, but whether any of them would yield power-law learning curves is difficult to answer. Even if the theories were stated in computational terms and allowed the collection of numerical results, there would still be the problems associated with discriminating power-law curves from exponential curves (Newell & Rosenbloom, 1981; Rosenbloom, 1986).

The closed-loop and open-loop distinction provides a better contrast between the theories. Adams and Pew's models cannot easily account for any open-loop behavior. Adams' memory trace could conceivably become sufficiently strong that simple movements could be performed in open-loop mode. Pew's schema instance can be forced into open-loop mode, since it is converted to a temporal sequence of muscle commands that theoretically could be executed entirely without feedback. Schmidt's theory is almost entirely open-loop, although it can give the appearance of closed-loop behavior. However, none of the theories give good explanations of how behavior could progress from closed-loop to open-loop as a result of practice.

Finally, only Schmidt's schema theory is able to explain the practice variability effect. Of course, this phenomenon was predicted by (and observed after) the introduction of his schema theory. As discussed by Schmidt (1975b), Adams' theory has no way to account for such a phenomenon. However, Frohlich and Elliott (1984) claim that even Schmidt's explanation is too weak and they present an alternative view on this subject. Although the empirical results are still inconclusive, it seems clear that at least in some cases the effect holds consistently. A solid theory of human motor learning should be able to account for at least some of these effects.

All of the theories (including Pew's with a hypothetical learning component) explain the psychological phenomena rather well (not surprisingly). However, they are all limited to simple, ballistic movements. Most work has been done on single-joint tasks in one dimension. Consequently the psychological theories have little to say about more complex tasks involving the interaction of multiple joints in non-trivial manners. As mentioned above, a computational model of these theories would facilitate a more thorough evaluation, and in general, could provide much needed insight to the nature of such proposed theories.

4. Computational Approaches to Motor Behavior

Now we consider models that have been implemented and that model jointed motor control by specifying the representation, performance, and learning processes as computational mechanisms. Again, we must choose some method or dimension to limit the systems we consider in this paper. In this case, we will focus on heuristic methods that employ learning techniques to get around weaknesses in computational power, along with systems that are heavily geared toward modeling some aspect of human motor control. This

means excluding much of the robotics literature in so far as the methods commonly used in that area are intended to find exact or optimal trajectories for mechanical manipulators. Also such methods tend to focus on low-level motor control, involving torques and voltages which we intend to ignore.

We will also exclude the literature on robot planning (e.g., Segre, 1987; Andrae, 1985), which is mainly concerned with problems of planning and operator sequencing, as opposed to the execution of varied limb movements. Of course, both this type of work and the low-level robotics work are important in their own right, but they are not directly related to the concerns of this paper. As we stated before, we are interested in theories or systems that address both the cognitive and physiological aspects of motor learning.

We start by considering several systems that have been designed as models of the human motor system or that have paid close attention to constraints imposed by this system. Then we turn to several other implementations that deal with the control of dynamic systems and that could conceivably be applied to jointed limbs, but which are not explicitly presented as models of human motor control. We close by examining the plausibility of both types of systems considered, with respect to the constraints and phenomena we introduced earlier.

4.1. Chunking Goal Hierarchies as a Model of Motor Learning

Rosenbloom (1986) presents a model that accounts for both the power law of practice and the reaction time data on stimulus compatibility. The latter phenomenon states that the reaction time to a given stimulus is inversely related to the extent to which that stimulus is compatible with the required response. For example, if a tone in the left ear requires a button press with the right hand, the reaction time will be longer than if a button press with the left hand were required.

Rosenbloom's architecture accounts for both of these phenomena. The representation consists of goal hierarchies that determine the solutions to particular tasks. These are mostly simple choice reaction-time tasks in which an appropriate response must be selected to a given stimulus. The nature of the goal hierarchies used to solve these tasks gives rise to the compatibility effect. Learning consists of creating chunks from sequences of subgoals that have been solved in a given situation, and the coinciding decrease of necessary processing explains the power-law of practice results.

This model can be viewed as an explanation of task-independent practice effects; however, we are specifically taking a motor learning perspective. It accounts for the two phenomena mentioned above, as well as a number of others, but it does not explain such phenomena as the speed-accuracy tradeoff, sequential dependencies, interference, discrimination, and reaction time distributions. The model has been applied only to tasks that involve minimal motor control – the execution of a selected response – and these responses have been modeled as primitive operators. However, one can imagine adapting the architecture to a level that would include lower-level motor primitives, allowing the creation of goal hierarchies of motor movements and subsequent chunking of portions of such hierarchies. A further limitation is the absence of a mechanism that can acquire the necessary goal hierarchies. Several extensions are described that could conceivably alleviate this limitation. Although Rosenbloom's theory is rather weak on issues of motor control, it is the only model we will consider that significantly address cognitive aspects. As such, it perhaps holds the greatest promise for addressing both high-level planning issues and low-level control issues. While there is some promise, the details have not been specified and so we turn to a model that focuses on low-level control issues.

4.2. Raibert's State-space Model of Motor Learning

Raibert's (1976) model of motor control and learning is one of the most serious attempts at carefully dealing with issues in the human motor system. He presents four properties of this system that he attempts to model: the ability to gain control of the limbs through experience, the ability to maintain control in the context of changes to the limbs, the ability to compensate for mechanical interactions between serial joints, and the ability to convert a desired movement from one representation to another. He qualifies this model as only a sub-system of a more complete model of motor control and learning. In particular, this sub-

system is responsible for acquiring appropriate feed-forward commands. This constraint allows the model to ignore interactions with the environment (which would require a feedback mechanism) and the issue of motor programs (although their existence is not questioned). The model is intended to process the class of ballistic movements, such as swatting a fly or swinging a bat.

Raibert's work focuses on the construction of a translator that takes descriptions of desired movements and converts them to commands directly interpretable by muscles or motors. The main difficulty of such a task is encoding or solving the mechanics of the particular limb. In Raibert's model, this information is extracted from the relationship between the limbs' inputs and outputs that result from previous attempts to move or position the limb. This extraction is made feasible by discretizing time and space. Time is sliced up into sufficiently small pieces to allow the simplification of the equations describing the motion of the jointed limb to a set of constants. These constants cannot be stored for the infinite number of possible states of the arm, so the state-space of the arm must be divided into regions or hyper-cubes. This state-space memory associates one set of constants with a given hyper-cube for the entire state-space. These constants are assumed to be satisfactory for "near" states, or ones within the same hyper-cube (given sufficiently small hyper-cubes). This process is referred to as a *piece-wise linearization* of the mechanical system representing the limb.

Learning in this model involves the storage of the parameters for individual states of the state-space memory. The constants stored are based on averages of previously calculated values for given situations. The calculation is based on the commands issued to the limb and the resulting accelerations (see Raibert, 1976, for details). As experience occurs, more parts of the state-space memory are visited and filled. On average, behavior will improve as a greater percentage of this memory is filled in. Noise in measuring the accelerations of the joints is dampened by averaging the calculated constants with existing values in a particular hyper-cube of the state-space memory. One might obtain practice variability effects from this model, since the novel task will be "closer" in the hyper-space to previous experience in the variable practice condition than in the constant practice condition.

4.3. Generalizing Motor Control Using Knowledge

One of the limitations of Raibert's (1976) tabular approach is that transfer between dissimilar movements is difficult or impossible. Atkeson (1987) presents an adaptive feed-forward method that overcomes this limitation. His system acquires a global model of the arm dynamics requiring the learning of only one set of parameters for the equations. This contrasts with the many sets of parameters necessary in tabular approaches, where each set of parameters applies only to the small, corresponding region of the state-space. Not only does Atkeson's approach reduce the number of necessary parameters, it also reduces the learning necessary to achieve a comparable level of performance. As stated above, the state-space method must "explore" the space of possible arm states and store parameters for each, whereas the global model can be learned in just a few "test movements". The system requires torque/force sensors at the wrist and arm joints in order to measure the torques resulting from the test movements. Given the relationships between the measured values and the commands, the system can infer a model of the rigid body dynamics for the arm. Note that the table lookup methods did not require torque sensing devices on the arm but only the ability to sense where the arm was currently positioned (in joint coordinates).

The global model allows the parameters to be used for controlling a variety of movements within the given arm's state space. Unfortunately, using the global model to assign the parameters introduces small errors, which arise because the arm is not entirely rigid, as the global model inference mechanism assumes. If the global model were modified to correct for these small errors in one particular trajectory, the performance on other movements would in turn deteriorate. Instead, Atkeson includes a mechanism for learning single trajectories that takes advantage of both the global model and the feedback information from a particular attempt at executing the trajectory. Given several practice attempts, the commands for the trajectory can be improved to a level arbitrarily close to the sensitivity of the manipulator hardware. The introduction of

single trajectory learning mechanism involves altering the control system memory to allow the storage of commands for particular trajectories. The details of this memory are not discussed, and it appears to be an unwieldy addition to the system.

For future research, Atkeson proposes the use of local models that would store the more correct dynamic model for local portions of the space. This proposal involves either learning the dynamics of a "central" movement for a set of similar movements, or a tabular approach giving the dynamics for a local portions of the space. Either way, the local model would serve as a correction factor to the global model when generating the feed-forward commands of a movement related to the local model. A unique feature of this proposal is that it effectively suggests a hierarchy of models. This allows a tradeoff to be made between the generality of the global models and the accuracy of the local models that would "gain the benefits of each and the drawbacks of none" (p. 30).

4.4. A Connectionist Approach to Hand-eye Coordination

Connectionist and neural net architectures have received considerable attention recently as models of human cognitive processes. Mel (1988) presents a robot arm controller called MURPHY that utilizes such an architectural framework. Although he did not specifically intend this system as a psychological model, the design process was constrained by knowledge of nervous system structures and their operation.

The architecture is based on two interconnected sets of neuron-like units. A visual array represents the field of view and a kinematic population represents the angles of the three joints that are controlled by MURPHY. These units are overlapping so that a single image or joint angle will activate a small population of units; this distinguishes the approach from state-space schemes. Learning involves the creation of weighted associations between these two populations of units. The visual units that are activated by the joints are associated with the joint angle units that describe the position of the arm. Because of the overlapping structure of these populations, the level of activation for a given set of units decays gradually as the arm moves away. Training consists of stepping through a representative portion of the possible joint configurations and creating the weighted associations.

After training, MURPHY can "grab" a visually presented object. The distance from the tip of the arm to the goal is evaluated and a move is selected that will reduce the distance by the greatest amount. This is described as an internal search, after which the arm is moved to the target destination in a single execution. Mel presents no results on learning, but it seems plausible that the number of search steps should decrease with the extent of training. Alternatively, the search trajectory should approach the straight line between the initial and target configurations as training is increased. The approach is an interesting one, although the current system is very limited in that there is no facility for the representation, execution, or acquisition of arbitrary arm trajectories. Still, it bears further attention as MURPHY continues to be developed.

4.5. Adaptive Feedback Control Models

All of the systems we have considered in this section have either used a constant feedback controller or ignored feedback all together. Improvements in performance were gained by modifying the commands responsible for generating the original movement. There is also considerable research in the area of adaptive mechanisms for feedback control; that is, feedback controllers that learn from errors in previous experience. Several of these studies have focused on the "pole balancing" task (Michie & Chambers, 1968), which consists of a cart on a one-dimensional track with a pole attached via a hinge. The cart can be moved left or right with a constant force. The goal is to keep the pole in a near vertical position by selecting appropriate sequences of left and right forces on the cart. Although these systems have not been proposed as models of human motor control, in some cases they have been associated with claims as to the viability of the approach for robotics in general (Sutton, 1984; Selfridge, Sutton, & Barto, 1985).

Michie and Chambers (1968) implemented a program, BOXES, utilizing a reinforcement learning mechanism in the pole balancing domain. They used an independent-association approach that involved discretiz-

ing the environment into a state space using pre-defined ranges. The average time to failure (falling of the pole) was updated from experience and the action with the longest average was selected for a given state. This should not be confused with Raibert's state-space memory, which discretized only memory and not experience. In BOXES, two cart-pole configurations are identical if they fall within the same region of the discretized space. That is, as the system learns the appropriate action to make in given states, the only generalization would be to other configurations considered as the same state. Sutton (1984) and Selfridge et al. (1985) present another reinforcement learning method using a linear-mapping approach. This also required the discretizing of the space into regions, but the choices made in a region are based on the probability of maintaining balance. The number of trials required to learn to balance the pole for some criterion number of time steps was significantly less than BOXES. Connell and Utgoff (1987) present another program, CART, that does not discretize the space and further reduces the required learning time. Their system employs a Shepard function to determine the degree of desirability of a particular state (cart-pole configuration), and learning involves adding a point from the cart-pole space with an evaluation of its desirability (provided by a critic) to the instance memory. CART learned to balance the pole in less than 16 trials, as opposed to an average of 75 for Selfridge et al. and 600 for BOXES.

Although these systems have no provision for motor programs or feed-forward control of any sort, they represent important progress in adaptive feedback control. A mechanism that can improve its responses to errors is an important part of a complete model of human motor control. However, the amount of increased understanding from these systems is limited. The approaches are made manageable by the simplicity of the pole balancing domain, in which there are only two operators. Also, when applied to the control of robotic arms, the complexity of the state space will increase dramatically. This does not mean that these problems cannot be overcome, but it does mean there remains a need for continued work in all areas of motor control.

5. Conclusions

In this paper, we have attempted to cover multiple facets of the literature on motor behavior and learning. This represents an enormous amount of previous work and some means of constraining the coverage must be employed. We have focused this survey around our goal of developing a computational theory of human motor behavior that can learn to perform complex tasks such as swinging a golf club, shooting a basket ball, or juggling. We selected some of the more significant phenomena as a basis for constraining the particular type of motor model we are interested in. The leading psychological theories were considered in this context, followed by a number of implemented computer models and systems.

Although the psychological theories accounted for the phenomena rather well, we were unsatisfied with the operability of these models. Even if the effort were made to implement these theories, they would still be limited to simple, ballistic movements. The computational approaches have for the most part focused on low-level issues of controlling the hardware. These contributions tell us little about how humans control their limbs or what types of behaviors one can expect from humans in particular situations. In summary, there remains a need for a computational model of human motor control and learning.

What we are really interested in is a computational model of human motor learning of reasonably complex tasks, such as throwing balls, swinging a golf club, or drawing shapes. We assume that such movements are controlled by a motor program or schema. This excludes inherently feedback-oriented tasks, such as pole balancing, that really have no pre-programmed component. We are ultimately interested in demonstrating movements that are heavily feedback related but that still include pre-programmed components, such as walking and juggling. By keeping an eye to the phenomena demonstrated in the laboratory, as well as the previous work done on psychological theories and computational models, we hope to move toward a comprehensive computational model of this sort.

References

- Adams, J. A. (1971). A closed-loop theory of motor learning. *Journal of Motor Behavior*, 3, 111-149.
- Andreae, P. M. (1985). *Justified generalization: Acquiring procedures from examples*. Doctoral dissertation, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA.
- Armstrong, T. R. (1970). *Training for the production of memorized movement patterns* (Technical Report No. 26). Ann Arbor: University of Michigan, Human Performance Center.
- Atkeson, C. G. (1987). *Roles of knowledge in motor learning*. Doctoral dissertation, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA.
- Bartlett, F. C. (1958). *Thinking: An experimental and social study*. London: Allen & Unwin.
- Connell, M. E. & Utgoff, P. E. (1987). Learning to control a dynamic physical system. In *Proceedings of the Sixth National Conference on Artificial Intelligence* (pp 456-460). Seattle, WA: Morgan Kaufmann.
- Crossman, E. R. F. W. (1959). A theory of the acquisition of speed-skill. *Ergonomics*, 2, 153-166.
- Ells, J. G. (1969). *Attentional requirements of movement control*. Doctoral dissertation, University of Oregon.
- Frohlich, D. M. & Elliott, J. M. (1984). The schematic representation of effector function underlying perceptual-motor skill. *Journal of Motor Behavior*, 16, 40-60.
- Langolf, G.D., Chaffin, D.B., & Foulke, J.A. (1976). An investigation of Fitts' Law using a wide range of movement amplitudes. *Journal of Motor Behavior*, 8, 113-128.
- Melville, D. S. (1976). *A test of motor schema theory: Performance of a rapid movement task in the absence of knowledge of results*. Doctoral dissertation, University of Iowa.
- McCracken, H. D. & Stelmach, G. E. (1977). A test of the schema theory of discrete motor learning. *Journal of Motor Behavior*, 9, 193-201.
- Mel, B. W. (1988). Building and using mental models in a sensory-motor domain: a connectionist approach. In *Proceedings of the Fifth International Conference on Machine Learning* (pp. 207-213). Ann Arbor, MI: Morgan Kaufmann.
- Michie, D. & Chambers, R. (1968). BOXES: An experiment in adaptive control, In E. Dale & D. Michie (Eds.), *Machine Intelligence 2*. Edinburgh: Oliver and Boyd.
- Newell, A. & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive Skills and Their Acquisition*. Hillsdale, NJ: Erlbaum.
- Pew, R. W. (1966). Acquisition of hierarchical control over the temporal organization of a skill. *Journal of Experimental Psychology*, 71, 764-771.
- Pew, R. W. (1974). Human perceptual-motor performance. In B. H. Kantowitz (Ed.), *Human Information Processing: Tutorials in Performance and Cognition*. New York: Erlbaum.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 90, 1-7.
- Raibert, Marc (1976). A state space model for sensorimotor control and learning (Technical Report AIM-351). Cambridge, MA: Massachusetts Institute of Technology, Artificial Intelligence Laboratory.
- Rosenbaum, David A. (1977). Selective adaptation of "Command Neurons" in the human motor system. *Neuropsychologia*, 15, 81-91.

- Rosenbloom, P. S. (1986). The chunking of goal hierarchies. In J. Laird, P. Rosenbloom, & A. Newell (Eds.), *Universal Subgoaling and Chunking*. Boston, MA: Kluwer Academic Publishers.
- Salmoni, A. G. (1973). *Attention demands of a linear motor response*. Unpublished manuscript, Motor Behavior Laboratory, University of Michigan.
- Schmidt, R. A. (1975a). *Motor Skills*. Harper & Row, Publishers, New York.
- Schmidt, R. A. (1975b). A schema theory of discrete motor skill learning. *Psychological Review*, 82, 4, 225-260.
- Schmidt, R. A. (1982). The schema concept. In J. A. S. Kelso (Ed.), *Human motor behavior: An introduction*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Schmidt, R. A. & White, J. L. (1972). Evidence for an error detection mechanism in motor skills: A test of Adams' closed-loop theory. *Journal of Motor Behavior*, 4, 143-154.
- Schmidt, R. A., Sherwood, D. E., Zelaznik, H. N., & Leikind, B. J. (1985). Speed-accuracy trade-offs in motor behavior: Theories of impulse variability. In H. Heuer, U. Kleinbeck, & K.-H. Schmidt (Eds.) *Motor Behavior: Programming, Control, and Acquisition*. Berlin, Germany: Springer-Verlag.
- Segre, A. M. (1987). *Explanation-based learning of generalized robot assembly plans* (Technical Report UILU-ENG-2208). Urbana, IL: University of Illinois at Urbana-Champaign, Department of Electrical and Computer Engineering.
- Selfridge, O. G., Sutton, R. S., & Barto, A. G. (1985). Training and tracking in robotics. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence* (pp. 670-672). Los Angeles, CA: Morgan Kaufmann.
- Snoddy, G. S. (1926). Learning and stability. *Journal of Applied Psychology*, 10, 1-36.
- Sutton, R. S. (1984). *Temporal credit assignment in reinforcement learning*. Doctoral dissertation, Department of Computer and Information Science, University of Massachusetts, Amherst, MA.
- Zelaznik, H. N. (1977). Transfer in rapid timing tasks: An examination of the role of variability in practice. In D. M. Landers & R. W. Christina (Eds.), *Psychology of motor behavior and sport* (Vol. 1). Champaign, IL: Human Kinetics.